

Logicism: Exact Philosophy, Linguistics, and Artificial Intelligence

Richmond H. Thomason
Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260
U.S.A.

August 4, 1995

Extended Abstract/Rough Draft of a Paper Submitted to
The 1995 Society for Exact Philosophy Meeting
Calgary, Alberta, May 25-28, 1994

*This is a working draft:
version of August 4, 1995.
The material is volatile; do not quote.
Comments welcome.*

VISIT...

LANZAROTE
Caliente.COM

Logicism: Exact Philosophy, Linguistics, Computer Science, and Artificial Intelligence

1. Introduction

Although logicism is usually associated with a relatively unfashionable approach to the foundations of mathematics, there is more logicism around nowadays than you'd think. Mathematics isn't the only domain in which a logical approach pays off, nor do logicist projects have to be planned on a grandly ambitious scale.

I've found it useful to regard the more formal trends in a number of areas in which logic is applied as variations on the theme of logicism, many of them local and tailored to specific domains. In this paper I will try to articulate this insight. I want to suggest that a tradition that is generally considered to have died out with Carnap has led to viable, rewarding avenues of research in many disciplines, many of which are intensively pursued today.

We get the familiar relationship between philosophical varieties of logicism and the more technical varieties that have arisen in scientific disciplines: the philosophical versions tend to be ambitious and foundational. The scientific varieties are more constrained; they tend to be associated with specific areas of inquiry or "domains," and are connected with technical problems in formalization and the mathematical side of logic.

Here is a general picture.

1. Philosophical varieties

- (i) Aristotle
- (ii) Leibniz
- (iii) Frege
- (iv) Carnap's attempts to provide a logical basis for the physical sciences.
- (v) There are some contemporary general approaches, of which the most popular is probably situation theory. I am not sure how to classify it. Few of the situation theorists are philosophers; probably it best viewed as a set of ideas for attacking various scientific problems in formalization.

2. Scientific varieties

- (i) Logicism in mathematical logic.
- (ii) Logicism in computer science. Feature structure logic as an example of a type of approach that has been very rewarding in computer science.
- (ii) Linguistic logicism. Attempts by linguists and logicians to develop a "natural language ontology" (and, presumably, a logical language that is related to this ontology by formally explicit rules) that would serve as a framework for natural language semantics.
- (iii) Attempts in Artificial Intelligence (AI) to formalize common sense knowledge.

2. Aristotelian versus Fregean logicism

Where X is domain or topic of inquiry, X logicism is the view that X should be presented as an explicit axiomatic theory from which the rest can be deduced by a logic. Aristotle clearly states logicism as a kind of ideal in the *Organon*. He also introduces the notion of a domain, and indicates that each domain will have its own appropriate principles. He often remarks that “exactness” may differ from domain to domain; for instance, we shouldn’t expect the same exactness in sublunary physics as in celestial physics.

But a logicist program doesn’t emerge in Aristotle’s work: largely, I think, because the theory of logic that he develops is much too weak for representing even the most exact Aristotelian science. Thus, he is unable to use the logic systematically in the scientific work; there is certainly never any attempt to verify that all the reasoning that is used in this work conforms to the logic. And he is not in a position to address the differences in logic that presumably would answer to the distinction between the inexact and exact sciences.¹

The following three paragraphs are quoted verbatim from [Thomason 91].

There is a moral here about logicism. X logicism imposes a program: the project of actually presenting X in the required form. But for the project to be feasible, we have to choose a logic that is adequate to the demands of the topic. If a logic must involve explicit formal patterns of valid reasoning, the central problem for X logicism is then to articulate formal patterns that will be adequate for formalizing X .

The fact that very little progress was made for over two millennia on a problem that can be made to seem urgent to anyone who has studied Aristotle indicates the difficulty of finding the right match of topic and formal principles of reasoning. Though some philosophers (Leibniz, for one) saw the problem clearly, the first instance of a full solution is Frege’s choice of mathematical analysis as the topic, and his development of the *Begriffsschrift* as the logical vehicle. It is a large part of Frege’s achievement to have discovered a choice that yields a logicist project that is neither impossible nor easy.

By allowing the underlying logic to be vague and inexplicit, it is easy to make any well-argued science seem logically correct, by pointing to principles and plausible inferences. Formalizing the logic keeps the logicist honest, and makes it much more difficult to show that a topic is formalizable—with a logic like the *Begriffsschrift*, a painstaking process of formalizing the relevant mathematical statements and proofs is needed to demonstrate that the knowledge can be expressed and the reasoning can be captured.

Frege’s logicism is associated with two programmatic features: (1) the use of a single, underlying logic, and (2) the idea that the program requires (at least in the case of mathematics) a definitional treatment of the domain, so that the inferences required by the logicist program are obtained by providing an appropriate definitional extension of the single logic appealed to in (1). Neither of these are a necessary part of a logicist program.

¹Aristotle’s position seems to have been that the logic of the inexact sciences is the same as that of the exact ones; but you can’t always rely on the conclusions in inexact domains. This position is no longer tenable.

I don't have much to say about the first claim. It is hard to make philosophical sense of the idea of multiple logics that aren't reconcilable in terms of a single logic. But the foundational paradoxes make it hard to see what such a single, universal logic would be like. For the logicist programs that have developed in the second half of this century, it seems to me that the more interesting research issues have to do with the formalization of various domains, rather than with finding a unified logical treatment that is appropriate for all domains.

3. Extensions to the empirical world

The project of extending Frege's achievement to the empirical sciences has not fared so well. Of course, the mathematical parts of sciences such as physics can be formalized in much the same way as mathematics. Though the metamathematical payoffs of formalization are most apparent in mathematics, they can occasionally be extended to other sciences.² But what of the empirical character of sciences like physics? One wants to relate the systems described by these sciences to observations.

Rudolph Carnap's *Aufbau*³ was an explicit and ambitious attempt to extend mathematics logicism to science logicism, by providing a basis for formalizing the empirical sciences. The *Aufbau* begins by postulating elementary units of subjective experience, and attempts to build the physical world from these primitives in a way that is modeled on the constructions used in Frege's mathematics logicism.

Carnap believed strongly in progress in philosophy through cooperative research. In this sense, and certainly compared with Frege's achievement, the *Aufbau* was a failure. Nelson Goodman, one of the few philosophers who attempted to build on the *Aufbau*, calls it "a crystallization of much that is widely regarded as worst in 20th century philosophy."⁴

After the *Aufbau*, the philosophical development of logicism becomes somewhat fragmented. The reason for this may have been a general recognition, in the relatively small community of philosophers who saw this as a strategically important line of research, that the underlying logic stood in need of fairly drastic revisions.⁵

This fragmentation emerges in Carnap's later work, as in the research of many other logically minded philosophers. Deciding after the *Aufbau* to take a more direct, high-level approach to the physical world, in which it was unnecessary to construct it from phenomenal primitives, Carnap noticed that many observation predicates, used not only in the sciences but in common sense, are "dispositional"—they express expectations about how things will behave under certain conditions. A malleable material will deform under relatively light pressure; a flammable material will burn when heated sufficiently. It is natural to use the word 'if' in defining such predicates; but the "material conditional" of Frege's logic gives incorrect results in formalizing such definitions. Much of [Carnap 36-37] is devoted to presenting and examining this problem.

Rather than devising an extension of Frege's logic capable of solving this problem, Carnap suggests dropping the requirement that these predicates should be explicated by definitions. This relaxation makes it harder to carry out the logicist program, because a nat-

²See [Montague 62].

³[Carnap 28].

⁴[Goodman 63], page 545.

⁵I can vouch for this as far as I am concerned.

ural way of formalizing dispositionals is forfeited. But it also postpones a difficult logical problem, which was not, I think, solved adequately even by later conditional logics in [Stalnaker & Thomason 70] and [Lewis 73]. Such theories do not capture the notion of normality that is built into dispositionals: a more accurate definition of ‘flammable’, for instance, is ‘what will *normally* burn when heated sufficiently’. Thus, logical constructions that deal with normality offer some hope of a solution to Carnap’s problem of defining dispositionals. Such constructions have only become available with the development of nonmonotonic logics.

4. Linguistic logicism

Though work in philosophical logic and its applications continues the logicist tradition to some extent, logicist projects are largely out of fashion in philosophy, and much of the work on projects of this sort is being carried on in other disciplines.

In linguistics, a clear logicist tradition emerged from the work of Richard Montague, who (building to a large extent on Carnap’s work in [Carnap 56]) developed a logic he presented as appropriate for *philosophy logicism*. Montague’s extreme logicist position is stated most clearly in a passage in [Montague 69].

It has for fifteen years been possible for at least one philosopher (myself) to maintain that philosophy, at least at this stage in history, has as its proper theoretical framework set theory with individuals and the possible addition of empirical predicates. ... [But] philosophy is always capable of enlarging itself; that is, by metamathematical or model-theoretical means—means available within set theory—one can “justify” a language or theory that transcends set theory, and then proceed to transact a new branch of philosophy within the new language. It is now time to take such a step and to lay the foundations of intensional languages.⁶

Montague’s motivation for expanding his logical framework is the need to relate empirical predicates like ‘red’ to their nominalizations, like ‘redness’. He argued that many such nominalizations denote properties, that terms like ‘event’, ‘obligation’, and ‘pain’ denote properties of properties, and that properties should be treated as functions taking possible worlds into extensions. The justification of this logical framework consists in its ability to formalize certain sentences in a way that allows their inferential relations with other sentences to be captured by the underlying logic.

Philosophers other than Montague—not only Frege, but Carnap in [Carnap 56] and Church in [Church 51]—had resorted informally to this methodology. But Montague was the first to see the task of *natural language logicism* as a formal challenge. By actually formalizing the syntax of a natural language, the relation between the natural language and the logical framework could be made explicit, and systematically tested for accuracy. Montague developed such formalizations of several ambitious fragments of English syntax in several papers, of which [Montague 73] was the most influential.

The impact of this work has been more extensive in linguistics than in philosophy.⁷ Formal theories of syntax were well developed in the early 70s, and linguists were used

⁶[Montague 74], pages 156–157.

⁷It is hard to explain the lack of philosophical interest in the project. Recent linguistic work in natural language metaphysics, of the sort described, for instance, in [Bach 89], is loosely connected to earlier at-

to using semantic arguments to support syntactic conclusions, but there was no theory of semantics to match the informal arguments. “Montague grammar” quickly became a paradigm for some linguists, and Montague’s ideas and methodology have influenced the semantic work of all the subsequent approaches that take formal theories seriously.

As practiced by linguistic semanticists, language logicism would attempt to formalize a logical theory capable of providing translations for natural language sentences so that sentences will entail one another if and only if the translation of the entailed sentence follows logically from the translation of the entailing sentence and a set of “meaning postulates” of the semantic theory. It is usually considered appropriate to provide a model-theoretic account of the primitives that appear in the meaning postulates.

This methodology gives rise naturally to the idea of “natural language metaphysics,” which tries to model the high-level knowledge that is involved in analyzing systematic relations between linguistic expressions. For instance, the pattern relating the transitive verb ‘bend’ to the adjective ‘bendable’ is a common one that is productive not only in English but in many languages. So a system for generating derived lexical meanings should include an operator **ABLE** that would take the meaning of ‘bend’ into the meaning of ‘bendable’.

To provide a theory of the system of lexical operators and to explain logical interactions (for instance, to derive the relationship between ‘bendable’ and ‘deformable’ from the relationship between ‘bend’ and ‘deform’), it is important to provide a model theory of the lexical operators. So, for instance, this approach to lexical semantics leads naturally to a model-theoretic investigation of ability,⁸ a project that is also suggested by a natural train of thought in logicist AI.⁹

Theories of natural language meaning that, like Montague’s, grew out of theories of mathematical language, are well suited to dealing with quantificational expressions, as in

4.1. *Every boy gave two books to some girl.*

In practice, despite the original motivation of his theory in the semantics of word formation, Montague devoted most of his attention to the problems of quantification, and its interaction with the intensional and higher-order apparatus of his logical framework.

But those who developed Montague’s framework soon turned their attention to these problems, and much of the later research in Montague semantics—especially David Dowty’s early work in [Dowty 79] and the work that derives from it—concentrates on semantic problems of word formation, which of course is an important part of lexical semantics.¹⁰

tempts to exploit language as a source of insights into the nature of distinctively human patterns of thought about what might be called the common sense world. I am thinking here of works like [Cassirer 55] and of [Jespersen 65]. Both of these projects grew out of a rich philosophical tradition: Cassirer’s work, in particular, is firmly rooted in the European Kantian tradition. And, of course, there has been much work in the phenomenological tradition—which, however, has been much less formal.

⁸That the core concept that needs to be clarified here is ability rather than the bare conditional ‘if’ is suggested by cases like ‘drinkable’. ‘This water is drinkable’ doesn’t mean ‘If you drink this water it will have been consumed’. (Of course, ability and the conditional are related in deep ways.) I will return briefly to the general problem of ability in Section 7.5, below.

⁹See, for example, [Thomas *et al.* 90].

¹⁰This emphasis on compositionality in the interpretation of lexical items is similar to the policy that Montague advocated in syntax, and it has a similar effect of shifting attention from representing the content of individual lexical items to operators on types of contents. But this research program seems to require a much deeper investigation of “natural language metaphysics” or “common sense knowledge” than the syntactic program, and one can hope that it will build bridges between the more or less pure logic with which Montague worked and a system that may be more genuinely helpful in applications that involve representation of and reasoning with linguistic meaning.

5. Case studies in linguistic logicism

I'll illustrate the use of nonmonotonic formalisms in natural language semantics with several case studies. In these studies, I'm merely trying to motivate the use of a nonmonotonic formalism in the semantics of words, and to suggest how it might be applied to some of the immediate problems that arise in this area. At the date of this version, I have not tried to work out the details. At this point, the abstract will become much more sketchy.

5.1. The '-able' suffix

The natural way to define ' x is water-soluble' is

5.1. If x were put in some water, then x would dissolve in the water.

So at first glance, it may seem that the resources for carrying out the definition that Carnap found problematic will be available in a logic with a subjunctive conditional. But suppose it so happens that if one were to put this salt in some water, it would be this water, and this water is saturated with salt. The fact that the salt would not then dissolve is no reason why the salt should count as not water-soluble. This and other such thought experiments indicate that what is wanted is not the bare subjunctive conditional, but a "conditional normality" of the sort that is used in some nonmonotonic formalisms.¹¹

In a circumscriptive framework, normality is obtained by conditions on a number of abnormality predicates, which are then circumscribed, or minimized relative to certain background assumptions, in obtaining models of the nonmonotonic theory. Events are an appropriate locus for organizing these abnormality predicates not only in the case of dissolving, but in many other cases of interest for purposes of lexical decomposition.

It is convenient to think of events as classified by a system of event types, from which abnormalities and other features are inherited. In treating the dissolving example, I will make the following assumptions.

- 5.i. There is an event type ϕ of *put-in events*.¹² Associated with this type (and, by inheritance, with events falling under it) there is a container $container(\phi)$ and a thing moved $movee(\phi)$.
- 5.ii. The event type ϕ has a subtype ϕ_1 , in which $container(\phi)$ is a quantity of water and $movee(\phi)$ is a quantity of salt. There is an abnormality predicate associated with ϕ_1 .
- 5.iii. There is an event type ψ of *dissolving events*. I assume that associated with this event type (and, by inheritance, with events falling under it) there is an inception, a body, and a culmination (where the first two are events and the last is a state); also, an associated medium $medium(\psi)$ and a thing dissolved $dissolvee(e)$; also an abnormality predicate.

It will follow from general considerations about the event type ψ that if a ψ -normal event of this type occurs, its associated culmination state will also occur. (See the remarks below on telicity.)

Given this information about event types, the sort of analysis that I currently favor for dissolving amounts to this.

¹¹See [Boutilier 92], [Asher & Morreau 91].

¹²This event type itself has a compositional analysis, but we can ignore that for purposes of the example.

- 5.iv. Every ϕ_1 -normal ϕ_1 -event e_1 is also the inception of a ϕ -event e_2 such that
 $container(e_1) = medium(e_2)$ and $movee(e_1) = dissolvee(e_2)$.

This analysis invokes notions that have come to light in accounting for other phenomena in the analysis of word meaning. I will pass directly to these other phenomena, but will return to the problem of dispositionals briefly later, when I discuss ability.

5.2. Telicity

I am abstracting here away from all problems having to do with time and the progressive, and concentrating on the relation between a telic event and its culminating state.¹³ The most important feature of the type of telic events is that these events have three associated parts: the inception, the body, and the culmination. The inception is an initiating event. The culmination is the state that normally results. (Since the beginning, the theory of planning has concentrated on features of culminations, since these represent properties of the state that can be assumed to result if the agent performs an action.) The body is the process that normally leads to the culmination; often (as in closing a door or filling a glass), the body will consist of stages in which the goal is progressively achieved.¹⁴ We can lay it down as a general default on telic events that the culmination of such an event will occur if its inception and body occurs. In many cases (like dissolving, or filling a glass from a tap, but not like filling a glass from a pitcher) the body will also normally occur if the inception occurs. Thus, I am likening unfulfilled telic events to Manx cats—they are objects that belong to a type that normally has a certain part, but that for some *ad hoc* reason happen to lack this part.

5.3. Agency

Some formalisms of agency in AI involve a separation of events into those that are in an agents' immediate repertoire and those that are not.¹⁵ If such a division is adopted for linguistic purposes, we can capture agency—at least, for telic events that normally follow from their inceptions—as follows.

- 5.v. $Do(x, e)$ holds iff the inception e' of e is identical to an immediate action e'' that is performed by x , provided that the body of e'' occurs.¹⁶

For example, it follows from this account that in case someone puts a piece of salt in water and it then dissolves in the ordinary way, then this person has also performed the action of dissolving it in water, assuming that putting the salt in water is an immediate action. Moreover, the action of putting the salt in water will be the inception of the dissolving event.

On this treatment, we dispense with an explicit use of any causal notions in the analysis of agency—though causal notions are certainly implicit if we believe that there is a connection between sequences of events conforming to patterns of normality and causal sequences. Since, despite the contribution of [Shoham 88], an explicit theory of common sense causality is not likely to be easy, I prefer such eliminative accounts. However, I'm not sure if explicit causality can be eliminated in general from the theory of agency.

¹³It should be clear, though, that I have in mind an account that would relate the truth of a progressive sentence to the occurrence of the body of an event.

¹⁴For ideas that are in some respects similar to these, see [Steedman & Moens 87].

¹⁵See especially [Moore 90].

¹⁶I want to say that the body of a telic event occurs even if it is partial or incomplete.

5.4. Causality

The notion of causality is usually left unanalyzed in linguistic treatments of lexical decomposition. But theories of nonmonotonic reasoning offer some hope of either providing an account of causality in terms of defaults governing sequences of events, or at least of providing systematic relations between such defaults and causality. For the most extended treatment, see [Shoham 88]. This work provides another, independent reason for incorporating a theory of default reasoning in an account of the compositional semantics of words.

5.5. Ability

‘This water is drinkable’ doesn’t mean

5.2. *An attempt to drink this water will normally culminate in its being drunk.*

Rather, the meaning is

5.3. *Normally, one can drink this water.*

This linguistic example illustrates the need for an account of practical ability. I don’t think that such an account can be given without an extended background theory of practical reasoning. For that reason, the account that I’ll sketch here may seem circular or trivial. The reason (I hope) is that the background hasn’t been filled in.

Let’s suppose that there is a propositional constant *practicalAbnormality* that is used in practical reasoning to reject alternatives because of utility considerations. That is, if a contemplated practical alternative is shown to lead (perhaps with the aid of defaults) to this constant, the alternative has thereby been shown to one that can be ruled out of consideration. A qualitative account of practical reasoning would have to relate this constant to desires and contingent circumstances.

The definition of practical ability would then be the following, where \Box represents temporal necessity.

5.vi. $can(\phi) \leftrightarrow \neg \Box[\phi \rightarrow practicalAbnormality]$

5.6. Artifacts

Many artifacts are defined in terms of their normal uses. This suggests decompositional analyses such as the following example.

5.vii. A *fastener* is an object x such that, where ϕ is the event type of using x , every ϕ -normal occurrence of an event e of type ϕ is such that *purpose*(e) is to fasten an object to another object.

6. Logicism in computer science

Note: This section is still tentative and in rough draft.

Computer science has raised the art of formalizing local domains to new levels of sophistication. Because of the training that computer scientists receive, this sort of work is a familiar and highly valued area of research, whether the formalization can be equipped with a compiler or otherwise implemented, or remains only an abstract specification of a problem.

Since mathematical logic has heavily influenced the thinking of computer scientists, the familiar elements of logical formalization are readily recognizable in the computational work;

syntactic specification of a language, characterization of inferential relations, elaboration of a model theoretic semantics, and use of the formalism for representing information in application domains. Naturally enough, developments in computer science have concentrated more attention on the algorithmic properties of formalisms, and much of the computational research is concerned with the algorithmic complexity of problems associated with various formalisms. The availability of computers and the needs of applications have also drawn attention to features like the maintainability of large systems, and the naturalness of representations, that previously were of little or no attention to logicians.

Logicist enterprises in computer science are usually associated with declarativism; the issues that emerged in debating the merits of procedural and declarative formalisms shed new light, it seems to me, on the value of logicist projects.

A number of logicist strategies have arisen in the computational arena that are of great practical importance, and that also are very interesting models of formalization. These strategies rely on limited formal techniques in order to obtain a proper balance between expressivity (the ability to represent enough information) and computational considerations such as implementability and efficiency. Some examples are: (1) logic programming, (2) unification formalisms in grammar, (3) taxonomic logics. If time permits, I'll develop one or more of these as case studies.

7. Common sense logicism

John McCarthy's logicist program in AI represents a version of common sense logicism that is similar in motivation to linguistic logicism, but that in many ways is more ambitious, and that has inspired a great deal of important work.

To a certain extent, the motives of the common sense logicism overlap with Carnap's motives for the *Aufbau*. The idea is that the theoretical component of science is only part of the overall scientific project, which involve situating science in the world of experience for purposes of testing and application; see [McCarthy 84] for explicit motivation of this sort, as well as [McCarthy 86] and [McCarthy 89].

I will describe the research issues in this area in my talk, but there is not much need to put much of this into a paper. I have already published a paper on McCarthy's logicism; see [Thomason 91]. And there are several extensive publications dealing with the formalization of common sense; see [Hobbs & Moore 90] and [Davis 90].

Bibliography

References

- [Asher & Morreau 91] Nicholas Asher and Michael Morreau. “Commonsense entailment: a modal theory of nonmonotonic reasoning.” In J. Mylopoulos and R. Reiter, eds., *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, Los Altos, California, 1991, pp. 387–392.
- [Bach 89] Emmon Bach. *Informal Lectures on Formal Semantics*. State University of New York Press, Albany, NY, 1989.
- [Boutilier 92] Craig Boutilier. *Conditional logics for default reasoning and belief revision*. Technical Report KRR–TR–92–1, Computer Science Department, University of Toronto, Toronto, Ontario, 1992.
- [Carnap 28] Rudolph Carnap. *Der logische Aufbau der Welt*. Weltkreis-Verlag, Berlin-Schlactensee, 1928.
- [Carnap 36-37] Rudolph Carnap. “Testability and meaning.” *Philosophy of Science* 3, pages 419–471, 1936, and *Philosophy of Science* 4, pages 1–40, 1937.
- [Carnap 56] Rudolph Carnap. *Meaning and necessity*, 2nd edition. Chicago University Press, Chicago, 1956.
- [Cassirer 55] Ernst Cassirer. *The Philosophy of Symbolic Forms*, volumes 1–3. Yale University Press, New Haven, CT, 1955. [These works were first published (in German) in 1923–29.]
- [Church 51] Alonzo Church. The Need for Abstract Entities in Semantic Analysis. *Proceedings of the American Academy of Arts and Sciences* 80, pages 100–112, 1951.
- [Davis 90] Ernest Davis. *Representations of commonsense knowledge*. Morgan Kaufmann, Los Altos, California, 1990.
- [Dowty 79] David Dowty. *Word meaning and montague grammar*. D. Reidel, Dordrecht, Holland, 1979.
- [Goodman 63] Nelson Goodman. “The significance of *Der logische Aufbau der Welt*. In P. Schilpp, ed., *The philosophy of Rudolph Carnap*, Open Court, LaSalle, Illinois, pp. 545–558, 1963.
- [Hobbs & Moore 90] Jerry Hobbs and Robert Moore. *Formal theories of the commonsense world*. Ablex, Norwood, New Jersey, 1988.
- [Jespersen 65] Otto Jespersen. *The philosophy of grammar*. W.W. Norton, New York, 1965. [This work was originally published in 1924.]
- [Lewis 73] David Lewis. *Counterfactuals*. Harvard University Press, Cambridge MA, 1973.

- [McCarthy 84] John McCarthy. “Some expert systems need common sense.” In H. Pagels, ed., *Computer culture: the scientific, intellectual and social impact of the computer*. Volume 426, Annals of the New York Academy of Sciences, pages 129–137, 1984.
- [McCarthy 86] John McCarthy. “Applications of circumscription to formalizing common sense.” *Artificial Intelligence*, 1986.
- [McCarthy 89] John McCarthy. “Artificial intelligence and logic.” In R. Thomason, ed., *Philosophical logic and artificial intelligence*, Kluwer Publishing Co., Dordrecht, Holland, pages 161–190, 1989.
- [Montague 62] Richard Montague. “Deterministic theories.” In *Decisions, Values, and Groups*, Volume 2. Pergamon Press, Oxford, pages 325–370, 1962. Reprinted in [Montague 74], pages 303–359.
- [Montague 69] Richard Montague. “On the nature of certain philosophical entities.” *The Monist* 53, pages 159–194, 1969. Reprinted in [Montague 74], pages 148–187.
- [Montague 73] Richard Montague. “The Proper treatment of quantification in ordinary English.” In J. Hintikka, J. Moravcsik, and P. Suppes (eds.), *Approaches to natural language: proceedings of the 1970 Stanford workshop on grammar and semantics*. D. Reidel, Dordrecht, Holland, pages 221–242, 1973. Reprinted in [Montague 74], pages 247–270.
- [Montague 74] Richard Montague. *Formal philosophy: selected papers of Richard Montague*. Edited by R. Thomason. Yale University Press, New Haven, CT, 1974.
- [Moore 90] Robert Moore. “A formal theory of knowledge and action.” In [Hobbs & Moore 90], pp. 319–358.
- [Shoham 88] Yoav Shoham. *Reasoning about change*. MIT Press, Cambridge, Massachusetts, 1988.
- [Stalnaker & Thomason 70] Robert Stalnaker and Richmond Thomason. “A semantic analysis of conditional logic.” *Theoria* 36, pages 23–42, 1970.
- [Steedman & Moens 87] Mark Steedman and Marc Moens. “Temporal ontology in natural language.” In *Proceedings of the 25th annual conference of the Association for Computational Linguistics*. Association for Computational Linguistics, 1987, pp. 1–7.
- [Thomas *et al.* 90] Becky Thomas, Yoav Shoham, and Anton Schwartz. “Modalities in agent-oriented programming.” Unpublished manuscript, Computer Science Department, Stanford University, Stanford, CA, 1990.
- [Thomason 91] Richmond Thomason. “Logicism, artificial intelligence, and common sense: John McCarthy’s program in philosophical perspective.” In *Artificial intelligence and mathematical theory of computation*, Vladimir Lifschitz, ed., Academic Press, San Diego, 1991, pp. 449–466.